

*Sub  
A  
B/C*

1. A method for scheduling prefetches into a cache of a data storage system, the method comprising:
  - 3      remotely modeling the dynamic operation of the cache;
  - 5      the remotely modeling including providing a model of data elements stored
  - 6      within the cache; and
  - 7      making a cache management decision based upon the model.
- 10     2. The method of claim 1, wherein making a cache management decision comprises:
  - 13      intercepting a request for a data element from a stream of Input/Output (I/O)
  - 14      data requests passed between a host and a storage device of the data storage system;
  - 15      and
  - 16      determining whether to schedule a prefetch of a data element logically
  - 17      successive to the requested data element in accordance with the contents of the
  - 18      cache as indicated by the model.
- 21     3. The method of claim 1, wherein the cache is a least recently used (LRU) cache.
- 24     4. The method of claim 2, wherein the LRU cache is a native LRU-only cache, and further comprising the step of leaving the native LRU-only cache substantially unmodified while conducting the steps of claim 2.

1       5. The method of claim 2, wherein determining whether to schedule a prefetch of  
2 data into the cache further comprises checking the model to determine whether the requested  
3 data element is likely to be present ~~within~~ in the cache.  
4

5       6. The method of claim 1, wherein determining whether to schedule a prefetch  
6 further comprises examining the history of a second data element stored logically adjacent to  
7 the requested data element in the storage device.  
8

9       10      7. The method of claim 1, wherein remotely modeling the cache further comprises:  
11            12      determining the size of the cache;  
13            14      periodically fetching the I/O rate of the cache; and  
15            16      periodically fetching the hit rate of the cache.  
17

18       8. The method of claim 1, wherein remotely modeling the cache further comprises  
19            20      periodically calculating a single reference residency time (SRRT) for a data element within  
21      the cache.  
22

23       9. The method of claim 1, wherein remotely modeling the cache further comprises  
24            25      the step of treating a requested data element as the youngest member of the cache when the  
26      requested data element is already present in the cache.  
27

- 1        10. The method of claim 1, wherein remotely modeling the cache further comprises  
2        determining whether the data element preceding the requested data element in a sequential  
3        stream of data is also present in the cache.  
4  
5
- 6        11. The method of claim 1, wherein remotely modeling the cache further comprises  
7        assigning a priority value to the requested data element.  
8  
9
- 10       12. The method of claim 11, wherein assigning a priority value further comprises  
11       assigning a priority value comprising the priority value assigned to the preceding data element  
12       plus one when the preceding data element is found to be present in the cache.  
13  
14
- 15       13. The method of claim 11, wherein determining whether to schedule a prefetch of  
16       a data element further comprises comparing the priority value of the requested element with  
17       a dynamic threshold.  
18  
19
- 20       14. The method of claim 13, further comprising prefetching the requested data  
21       element into the cache if the priority value of the requested data element is greater than the  
22       dynamic threshold.  
23  
24
- 25       15. The prefetch method of claim 1, further comprising periodically reevaluating the  
26       performance of the cache model.

1           16. The method of claim 15, wherein periodically reevaluating the performance of  
2           the cache further comprises determining if the dynamic threshold used in the internal model  
3           of the cache accurately models the performance of the cache.  
4  
5

6           17. The method of claim 16, wherein determining if the dynamic threshold  
7           accurately models the performance of the cache comprises comparing the performance of the  
8           dynamic threshold with an alternate dynamic threshold.  
9  
10

11           18. The method of claim 15, further comprising automatically updating the dynamic  
12           threshold used in the internal model of the cache when another dynamic threshold is deemed  
13           to be more effective.  
14  
15

16           19. The method of claim 1, wherein making a cache management decision  
17           comprises deciding to schedule a prefetch, and further comprising scheduling a prefetch by  
18           sending an I/O request to the cache.  
19  
20

21           20. A method for scheduling prefetches in a data storage system having a host and  
22           a cache, the method comprising the steps of:  
23  
24           providing a cache for caching Input/Output (I/O) data;  
25  
26           providing a prefetch module remote to the cache;

Sub  
A4

remotely modeling the cache within the prefetch module and determining whether to schedule a prefetch of data into the cache according to the results of the step of remotely modeling the cache, the step of remotely modeling the cache module further comprising:

examining the history of a data element in the cache;

assigning a priority value to the data element according to its history;

comparing that priority value to a predetermined threshold value;

determining the size of memory used in the cache;

periodically fetching the I/O rate of the cache from the cache;

periodically fetching the hit rate of the cache from the cache; and

determining a single reference residency time for a data element within the cache;

intercepting a stream of I/O information from the host to the cache to locate a requested data element;

determining if the requested data element in the stream of I/O information is already present within the cache;

making the requested data element the youngest member of the cache;

determining if the data element preceding the requested data element is present in the cache;

assigning a priority value to the requested data element;

periodically reevaluating the performance of the cache versus an internal model of the cache if the number of I/O requests received by the cache is greater than a predetermined number;

Sub  
A4

- updating the dynamic threshold used in the internal model of the cache if the dynamic threshold value does not adequately model the performance of the cache;
- comparing the priority value of the requested data element with the dynamic threshold value; and
- prefetching the requested data element if the priority value of the requested data element is greater than the dynamic threshold value by passing an I/O request of the data element to the cache.

21. A data prefetch scheduling system comprising:

- a cache configured to communicate with a host; and
- a remote prefetch module configured to communicate with the host and the cache and configured to determine whether to schedule a prefetch of data into the cache; and
- a modeling module operating within the prefetch scheduling module configured to model the cache.

22. The data prefetch scheduling system of claim 21, wherein the cache comprises a least recently used (LRU) cache.

23. The data prefetch scheduling system of claim 22, wherein the LRU cache is a native LRU-only cache that is not internally modified.

1                   24. The data prefetch scheduling system of claim 21, wherein the remote prefetch  
2 module further comprises a calculation module configured to compare a priority value  
3 assigned to a data element to a threshold value and determine whether to schedule a prefetch  
4 of the data element.

5

6

7                   25. The data prefetch scheduling system of claim 21, wherein the remote prefetch  
8 module further comprises a dynamic threshold optimization configured to calculate and  
9 update a dynamic threshold used in determining whether to prefetch data.

10

11

12                  26. The data prefetch scheduling system of claim 21, wherein the remote prefetch  
13 module is configured to model the cache for use in determining when to prefetch I/O data into  
14 the cache.

15

16

17                  27. The data prefetch scheduling system of claim 21, wherein the remote prefetch  
18 module is configured to prefetch data into the cache according to a priority scheme that takes  
19 into account the run length of each sequential I/O stream.

20

21

22

23                  28. The data prefetch scheduling system of claim 21, further comprising a prefetch  
24 request module, the prefetch request module configured to request a data I/O from the cache  
25 when the remote prefetch module determines that a prefetch is to be conducted.

26

Add  
A5